

1

Motivation, Strategy, and Definition

1.1 Closure as Axiom

Knowledge closure is, roughly, the principle that any agent who knows P and recognizes that P implies Q knows – or is in a position to know – Q.¹ Although the principle has been challenged – most famously by Robert Nozick and Fred Dretske – most contemporary epistemologists remain committed to it.² For many, indeed, that commitment is so firm that they view a theory that rejects closure to be seriously undermined if not refuted outright for that reason alone.³

A set S is closed under an operation O when applying O to members of S delivers members of S. The set of natural numbers, for example, is closed under addition: adding two natural numbers always produces a natural number. This follows from the Peano axioms for the natural numbers and the definition of addition. The set of one's ancestors is closed under the parent-relation: any

¹ Refinements will come in §§1.4–1.7.

² The *loci classici* of closure denial are Dretske 1970 and Nozick 1981, chapter 3.

³ “[T]he idea that no version of [closure] is true strikes me, and many other philosophers, as one of the least plausible ideas to come down the philosophical pike in recent years.” (Feldman 1995, 487) “[I]f a philosopher advances a view that forces us to reject closure, that should be taken as a *reductio* of that philosopher’s view.” (Fumerton 1995, 131) “Robert Nozick’s counterfactual analysis of knowledge is famously inconsistent with intuitive closure, but that is usually taken as a reason for rejecting the analysis, not for rejecting closure.” (Williamson 2000b, 117)

parent of one's ancestor is also one's ancestor. This is explicable by appeal to the fact that "x is an ancestor of y" means "y is a direct or indirect descendant of x." And the set of true propositions is closed under deductive consequence: any consequence of a true proposition is also a true proposition. This is explicable by appeal to the nature of deductive consequence and a soundness proof. Closure principles typically have some sort of explanatory ground.⁴

However, this does not seem to be the case for knowledge closure. There is no argument of the form:

- (1.) Knowledge is like so ...
- (2.) Inference is like so ...
- (3.) Therefore, knowledge closure is true,

where the premises cite uncontroversial characteristics of knowledge and inference.⁵

Nor does every serious epistemological view imply closure. While the closure-denying sensitivity accounts of Nozick and Dretske face numerous objections, they nevertheless exert considerable intuitive pull. A belief is sensitive when, were the belief false, the agent would not believe it.⁶ Sensitivity

⁴ Although not always. The Peano axioms themselves include two closure principles.

⁵ I will discuss a possible exception in §1.8 and §7.2.

⁶ This is the simplest version of sensitivity. Actual accounts are more elaborate. Some relativize to method: if the belief were false and the agent were to employ the same method, then the agent would not believe it by that method (Nozick 1981). (Nozick also includes an "adherence" condition: were the belief true and the agent to employ the same method, the agent would believe it by that method.) Others attach the modality to the agent's reason for belief rather than the belief itself: if the belief were false, the agent would not have the reason she has for believing it (Dretske 1971). Finally, others couch sensitivity in

is not closed under deductive entailment.⁷ Nevertheless, there is undoubtedly *something* unattractive about a belief's counting as knowledge when the agent would still have believed it if it were false. Perhaps this intuition is illusory in some way.⁸ But the intuition is there, and strong enough to motivate a variety of successor views that attempt to reconcile sensitivity with closure.⁹

So closure advocates cannot claim that every intuitively plausible epistemological view implies closure. But there is also no argument from relatively uncontroversial truisms about knowledge and deductive inference on the table. This isn't to say that there are no arguments for closure at all; far from it. But they typically proceed by presenting considerations, not directly in favor of closure, but rather against its denial. Those arguments deserve serious

probabilistic rather than modal terms: the probability that the belief is true given that the agent believes it is 1 (Dretske 1983) or high but not necessarily 1 (Roush 2005, who also includes a probabilistic version of Nozick's adherence condition). And there are many other versions. Nozick called sensitivity "variation," and called the conjunction of variation and adherence "sensitivity" (or "tracking"). Nevertheless, "sensitivity" is reserved for variation in the subsequent literature, a terminological tradition that I follow here.

⁷ "I have hands" is sensitive: if I were not to have hands – because of an unfortunate accident with a table saw, for example – then I would not believe that I do. But "I am not a handless brain in a vat (BIV) stimulated to have the very experiences I do have" is not sensitive: if I were a BIV, I would still believe that I am not a BIV (since my reasons for believing this, whatever they might be, would remain). Nevertheless, "I am not a handless BIV" follows from "I have hands."

⁸ Sosa 1999a suggests that the intuition results from a confusion of sensitivity with its contrapositive, safety. These are not equivalent since both incorporate subjunctive conditionals, which are not truth-preserving under contraposition.

⁹ See, for example, DeRose 1995, 1996, and 2010; Roush 2005 and 2012; Baumann 2012; Black 2008a; and Murphy & Black 2012.

attention.¹⁰ But it is *prima facie* surprising that this widely endorsed principle isn't derivable from undisputed characteristics of knowledge and inference.

Closure is instead typically treated as an independent epistemological axiom.¹¹ As such, it is thought to be warranted in the way that axioms often are: it is intuitively obviously correct.¹² But if closure is a primitive epistemological axiom, it's an unusually complex one. Like the parallel postulate in Euclidean geometry, whose complexity motivated attempts to derive it from the other simpler axioms, it stands out among seeming truisms concerning knowledge ("what you know you believe," "what you know is true," "what you know can't be accidentally true," and so on) as begging for derivation from simpler axioms concerning knowledge and deductive inference.

The parallel postulate proved not to be so derivable. It also proved to be eliminable in favor of alternatives, giving rise to non-Euclidean geometries. Euclidean geometry is both intuitive and reasonably accurate as a representation of local observable space. But the parallel postulate turned out to be a dispensable theoretical posit rather than an unassailable geometric primitive in the representation of the geometry of the physical universe overall.

¹⁰ See [Chapter 10](#).

¹¹ "This principle seems to me something like an axiom about knowledge." ([Cohen 2002](#), 312).

¹² "That something like [closure] is true, I will be taking as a primitive epistemic fact. I'm unable to formulate an argument that [closure] is true, just as I cannot provide an argument that killing innocent children without cause is morally wrong. But just as I nevertheless take it to be obviously true that we shouldn't kill innocent children without cause in spite of my inability to argue for this truth, so I will be taking the truth of [closure]." ([Dodd 2012](#), 341).

I suggest that a similar situation holds with respect to closure. Closure is an expression of an undeniable truth: deductive inference is an excellent way to extend one's knowledge. But that undeniable truth is compatible with closure's strict falsehood as a universal feature of epistemic space; excellence is not undermined by failure under exceptional circumstances. Just as caution must be exercised when extrapolating from the apparent geometric characteristics of our local space to the universe overall, so caution must be exercised when extrapolating from the undeniable truth that closure reflects to its supposed status as a universal epistemic truth. Closure is also a dispensable theoretical posit rather than an immobile pivot around which the epistemological landscape must turn.

1.2 Why Care?

Many riches flow from the renunciation of closure. A prominent skeptical threat – that one can't know, for example, that one has hands unless one knows that one is not a handless brain in a vat – is defused in a way that respects our intuition that we don't know that skeptical hypotheses are false while preventing the spread of that ignorance to more pedestrian knowledge claims.¹³ There is no need to countenance highly unintuitive “easy knowledge” inferences. A plausible solution to the problem of “bootstrapping” becomes available.¹⁴ And there is no need to resort to various theoretical, semantic, or pragmatic maneuvers in order

¹³ See [Chapter 11](#). I examine the skeptical closure argument itself in [Chapter 3](#).

¹⁴ See [Chapter 11](#) for discussion of bootstrapping and easy knowledge.

to articulate views that incorporate closure, while at the same time conceding that most, if not all, of our knowledge is acquired from fallible sources.¹⁵

But fundamentally at stake, for me at any rate, is an untenable conception of the demands that an agent must satisfy in order to know. One knows by courtesy of internal and external conducive circumstances: you don't know where your car is by seeing it in the parking lot unless you remember what your car looks like, there's adequate lighting, light travels in a straight line, and so on. Call these *enabling conditions*. Such conditions must be in place for knowledge acquisition. But must the agent also know that they are in place?

It's hard to see why. There's no reason in general why an agent S's standing in a particular relation R to some fact, which requires that condition C is realized, requires that she also stand in R to C itself. My successfully maneuvering a car through an obstacle course requires that the brake pedal be appropriately connected to the brakes. But it doesn't require that I connected them. Why should S's knowing that P, which requires that enabling condition C is satisfied, require also that she know that C is satisfied?

Some – but not all – of the enabling conditions for S's knowledge of P are implied by P itself. Suppose, for example, that P is “the gas tank is empty,” which S believes as a result of consulting the gas gauge whose needle points at “E.” An enabling condition of S's knowing that the tank is empty this way is that the needle isn't stuck on “E.” If it were stuck on “E,” it would be so either while the

¹⁵ These include externalist accounts of both evidence and method, contextualism, pragmatic encroachment views, and safety accounts, among many others. It also includes brute-force reconciliations of closure with views that are not, on their face, closure-friendly by simply appending a closure principle; Sherrilyn Roush's 2005 tracking-with-closure account is an example.

tank isn't empty or while it is, coincidentally, empty. "The tank is empty" implies that the former possibility is not realized: if the tank is empty, then it's not the case that the tank isn't empty while the needle is stuck. Since it's an enabling condition of S's knowledge that the tank is empty that the needle isn't stuck (whether or not the tank is empty), it's also an enabling condition of that knowledge that the needle isn't stuck while the tank isn't empty.

This generalizes. For any enabling condition C, since knowledge of P is not compatible with the failure of C, it is also not compatible with the failure of C while some other fact is true, including $\sim P$. So $\sim(\sim C \ \& \ \sim P)$ is also an enabling condition, one that is implied by P. $(\sim C \ \& \ \sim P)$ is incompatible with knowledge of P for two reasons: it is incompatible with P itself – and so with the facticity of knowledge – and it is incompatible with C, a condition of S's knowledge of P given how she acquires that knowledge.

Finding closure intuitive, one could insist that S need only know that those enabling conditions that do follow from P are satisfied. So she needs to know that it's not the case that the needle is stuck while the tank isn't empty (if, at least, she recognizes the inferential relation), but she doesn't need to know that it's not the case that the needle is stuck while the tank is empty (and so also doesn't need to know that the needle isn't stuck *simpliciter*).

But this is intuitively arbitrary, with respect to both what S needs to know and what she is in a position to know. It's *prima facie* unintuitive that she needs to know that it's not the case that the needle is stuck while the tank isn't empty but not that the needle isn't stuck. And it's similarly unintuitive that she could be

in a position to know that it's not the case that the needle is stuck while the tank isn't empty but not in a position to know that the needle isn't stuck.

But if she does need to know that the needle isn't stuck then skepticism seems the inevitable result. The same considerations apply to any enabling condition.¹⁶ So S needs to know that each such condition is satisfied. That knowledge will, moreover, have its own enabling conditions. So S must know that those conditions are satisfied as well, and that each condition of *that* knowledge is satisfied, and so on. The imposition of such requirements seems destined for skepticism.¹⁷

1.3 Strategy

It's hard to argue against a principle that is widely treated as an epistemological axiom grounded in intuition. Even if a closure denier were to develop a theoretical account that rejects closure and that successfully answers every other conceivable objection (unlike those of Nozick and Dretske), she remains susceptible to the critique that her view doesn't preserve closure. She would likely be accused of acquiring the fruits of her theory by theft over honest toil: *of course* she has an answer to the skeptic, for example, but only because she hasn't

¹⁶ That is, for any such condition C, there is a condition $\sim(\sim C \ \& \ \sim P)$ that follows from P, and it seems correspondingly arbitrary that S needs to, and can, know that without also knowing $\sim(\sim C \ \& \ P)$ and so, simply, C.

¹⁷ Indeed, as we'll see in [Chapter 6](#), skepticism results from the demand that S need only know that those conditions hold that follow from P. (At this point I'm only describing a motivation for resisting closure, not an argument for doing so. The argument comes in the rest of this book.)

done the necessary hard work of providing such an answer that is compatible with closure.

In order to proceed under such dialectical circumstances, the arguments I will offer against closure are, I suggest, not just intuitively compelling, but are so from a variety of epistemological standpoints. I will also show that the intuition behind closure is not as forceful as it seems at first glance, and that ultimately it does not support closure. Finally, I will show that the abominable conjunction and spreading problems directed against closure denial can be answered.

The closure denier does owe an account of when, and why, closure fails. Such an account might attempt to isolate all and only closure failures; one would certainly expect this of a full-blooded, closure-denying theory of knowledge.¹⁸ A more modest aim is to identify conditions under which closure fails, without claiming that these are the only such conditions. The result might be less satisfying than a “closure fails if and only if X” account. But, for the closure denier’s purposes, no more than the “if” direction is required: it would suffice, and is more secure, to present a sufficient-but-perhaps-not-necessary account with broad theoretical and intuitive appeal. Developing a view satisfying the

¹⁸ This implies that some restricted version of closure is true. Since closure is surely not a property randomly distributed over inferences, there is some general characterization of those cases in which it does hold. Such a principle would, however, apply in a more restricted class of cases than would closure principles that are typically endorsed by those who identify themselves as closure advocates, and so will still count as closure denial in the relevant sense. (However, some philosophers who so identify themselves offer versions of closure that are in fact more restricted than those typically endorsed by mainstream closure advocates; [Baumann 2012](#) and [Roush 2012](#) are examples. It is, as a result, disputable whether they really should count as closure advocates.)

“only if” direction as well is likely to require a full-blooded theory of knowledge, in which case it will then run into the theft-over-honest-toil objection.

I will not, therefore, attempt to derive closure failure from some particular theory of knowledge (or class of such theories). This might disappoint some closure advocates since many of the objections against closure have in fact been directed against particular theories – especially those of Dretske and Nozick – that imply closure failure.¹⁹ But, on the face of it, such an argumentative strategy is inadequate; that T, which entails $\sim C$, is false does not imply that C is true.

So I will claim that there are conditions such that, when they are realized, closure fails, although there may well be other conditions with the same effect. Moreover, the cost of endorsing closure under those conditions will be very high indeed. The overall result will be that, far from closure denial’s being a theoretical disadvantage, it is incumbent on any defensible theory of knowledge that it accommodates closure failure.

In the remainder of this chapter I take up the challenging task of formulating a defensible closure principle. The next chapter presents a version of Dretske’s argument by counterexample, which appeals to putative counterexamples to closure. That argument will then structure the discussion for [Chapters 3–9](#), in which I examine the different strategies the closure advocate

¹⁹ Much of Hawthorne’s 2005 defense of closure, for example, is less an attack on closure denial *per se* than the presentation of counterexamples to Dretske’s conclusive-reasons account of knowledge (naturally enough, since he was responding to Dretske).

might adopt in way of responding to Dretske's argument.²⁰ The conclusion of [Chapter 9](#) is that each such strategy fails, and so Dretske's argument succeeds.

In [Chapter 10](#) I will respond to two popular arguments against closure denial: the abominable conjunction problem and the spreading problem. In the course of doing so I'll also examine closure-preserving contextualism and the non-skeptical invariantist closure denier's response to skepticism.²¹ In [Chapter 11](#) I'll examine the bootstrapping problem, epistemic circularity, and the relationship between knowledge and justification closure.

1.4 Defining Closure

Notwithstanding broad agreement that some sort of closure principle is true, it turns out to be very difficult to formulate a principle that is immune to counterexamples that are recognized as such by both friends and foes of closure. I will hereafter limit attention to single-premise closure, which concerns only inferences with one premise. As is well known, there are objections to multi-premise closure that don't apply to single-premise closure (but not vice versa). So, if single-premise closure is undermined, then so is closure overall.²²

²⁰ I provide a more detailed description of those chapters at the end of [Chapter 2](#), after Dretske's argument is in place.

²¹ Contextualists claim that the semantic value of "know" varies across contexts of knowledge attribution. Invariantists deny that there is such variation.

²² For similar reasons, I will not consider closure over inductive inferences. The problem for multi-premise closure is that small probabilities of error for each premise can add up so that, while each premise is probable, the conclusion is not. On probabilist conceptions of knowledge, according to which knowing P requires that P is probable on one's evidence, closure can fail as a result. This does not apply to single-premise inference; if P implies Q, the

The simplest version of closure – appearing primarily in studies of epistemic logic – is that, if S knows that P and P implies Q, then S knows that Q. But this is an obviously inadequate description of actual epistemic agents. If such an agent has no grip whatsoever on the fact that P implies Q, it is highly implausible that she nevertheless must know that Q.

A common formulation declares that if S knows *both* that P and that P implies Q, then S knows that Q (call this the *Classical Formulation*). But it doesn't follow from the antecedent that S even believes Q; knowledge of Q, however, requires belief that Q. And, even if S believes Q, it is compatible with this formulation that she doesn't do so because it follows from P. She could believe Q solely on the basis of wishful thinking and so, presumably, would not know it.

A more recent, and widely adopted, formulation is offered by John Hawthorne, inspired by Timothy Williamson's suggestion that closure is an expression of the capacity of deductive inference to increase what one knows.²³ "Williamson has an insightful take on the root of epistemic closure intuitions," says Hawthorne, "namely the idea that 'deduction is a way of extending one's knowledge'."²⁴ Call this *Williamson's insight*. Here is Hawthorne's formulation, with its scope and necessity made explicit and the clauses labeled for convenience:

probability of Q is at least as high as P. (Nevertheless, [Lasonen-Aarnio 2008](#) argues that the same problem can be extended to single-premise inferences.)

²³ [Hawthorne 2004 and 2005](#).

²⁴ [Hawthorne 2005](#), 41, fn. 6, quoting [Williamson 2000b](#), 117.

Hawthorne's Formulation

Necessarily, for all agents S and propositions P and Q: if (a) S knows that P and (b) competently deduces Q from P, thereby (c) coming to believe Q, while (d) retaining her knowledge that P throughout, then (e) she knows that Q.²⁵

Clause (d) is designed to exclude cases wherein, during the course of performing the inference, the agent somehow loses her knowledge of P (because, perhaps, the performance somehow brings misleading evidence to light). Note that closure, so formulated, is diachronic: clause (a) refers to S's knowledge at one time and clause (e) refers to her knowledge at a subsequent time. S's performance of the inference takes up the intervening time.²⁶

1.5 KC

Clause (b) of Hawthorne's Formulation replaces "knows that P implies Q" in the Classical Formulation. The extent of its departure from that formulation depends on how "competent deduction" is to be interpreted.

²⁵ Hawthorne 2004, 34. In Hawthorne 2005, 29 he substituted "comes to know that Q" for (e). However, and as Hawthorne recognized in the earlier work, an agent could know Q already, before performing the inference, and so satisfy the antecedent without satisfying the consequent (Hawthorne 2004, 34, fn. 86). But such a case should obviously not count as a counterexample to closure. So I cite his earlier formulation here.

²⁶ For this reason, Hawthorne's Formulation is not, strictly speaking, a closure principle, since closure principles specify conditions on set membership (at a time). For reasons that will soon be apparent, I won't attempt to revise it further.

A competent deduction might consist in only a single inferential step from P to Q or a sequence of such steps from P to Q. It is, however, better to characterize the latter as involving successive instantiations of closure rather than a single instantiation. After all, an intermediate conclusion in the sequence is the next inference's premise. If S doesn't know that intermediate conclusion, then she doesn't know the premise of the next inference. But, if so, it is unintuitive that she, nevertheless, must know the subsequent conclusion inferred from that premise. So knowledge of the ultimate conclusion requires that closure succeeds for each inferential step. We might as well, then, construe closure as applying to single-step inferences from the outset.

But a single-step inference seems to involve no more than the recognition that P implies Q, which is synchronic: one recognizes that P implies Q at a time, rather than across an interval of time. One might think that some span of time is involved, since S initially believes P and then acquires her belief in Q by performance of the deduction. But "S competently deduces Q from P" should not be understood to imply that S believes either P or Q. The former is the purpose, in part, of clause (a), and the latter of clause (c). Appeal is also made to closure in way of explaining "retraction" phenomena, wherein S, taking herself to not know Q and realizing that Q follows from P, proceeds to deny that she knows P, despite having previously claimed knowledge of P. S's competent, single-step deduction, then, just consists in her recognition that P implies Q, without commitment to either.

To recognize that P implies Q is not merely to know that it does (and so this does not amount to a reversion to the Classical Formulation). S might know

that P implies Q by testimony from a logician, without having any grip on the inferential relation herself. Some might think that this would suffice. But suppose S knows P, and knows that P implies Q by testimony. How do these pieces of knowledge fit together in order to deliver her knowledge of Q? Presumably by a *modus ponens* inference: she recognizes that, since P is true and P implies Q, Q is true as well. But perhaps she only knows that (P and (P implies Q)) implies Q by testimony as well. Then how do *these* pieces of knowledge fit together in order to deliver her knowledge of Q? Presumably by another MP inference from (P and (P implies Q)) and (if (P and (P implies Q)) then Q) to Q. But perhaps she only knows this by testimony as well ...

If we model S's relation to the fact that P implies Q as merely something S knows, and so just another proposition that she believes and can wield as a premise then, like Lewis Carroll's tortoise, she will never be in a position to detach the conclusion.²⁷ S's recognition that P implies Q cannot be construed as merely something that S knows and so believes, so that an independent disposition to infer Q from her beliefs that P and that P implies Q must be postulated. Rather, to recognize that P implies Q is to be *inherently* disposed to believe Q if one believes P, and not-P if one believes not-Q; if S does not have

²⁷ [Carroll 1895](#). Closure is sometimes represented as involving a *modus ponens* inference from S's beliefs that P and that P implies Q to Q. As per Carroll's story, however, that's a mistake. S infers from P to Q, not from P and (P implies Q) to Q.

those dispositions then S does not recognize that P implies Q.²⁸ That disposition is then manifested by S's believing Q when she believes P.²⁹

This suggests the following definition of closure:

Knowledge Closure (KC)

Necessarily, for every agent S and propositions P and Q: if

(a) S knows that P while (b) recognizing that P implies Q, and (c)

believes Q on the basis of her belief that P and recognition that P

implies Q, then (d) S knows that Q.³⁰

²⁸ Lasonen-Aarnio 2008 makes essentially the same point against knowledge-of-inference formulations of closure, albeit in the course of arguing for a competent-deduction formulation of closure.

²⁹ Does recognition imply knowledge? Perhaps not. Suppose that S receives excellent, though misleading, testimonial evidence to the effect that P doesn't imply Q. On some views, excellent but misleading evidence that R is false destroys knowledge of R, so that S no longer knows that P implies Q. But she still seems to recognize that it does, even while having evidence to the effect that it doesn't. If so, recognition does not imply knowledge. However, it might reasonably be thought that such misleading evidence also undermines her acquisition of knowledge of Q in virtue of that recognition. If so, recognition (as cited in the closure principle) would require knowledge of the implication relation, and thereby rule out the existence of such undermining evidence.

³⁰ The original clause (d) – that S retains her knowledge of P throughout – is no longer required; its purpose – ensuring that S knows P while competently inferring – is achieved by the word “while” in clause (a) of KC. I have also substituted Hawthorne's “coming to believe Q” with clause (c). Suppose that S believes Q before recognizing that it follows from P but does not (yet) know it. (Perhaps she believes it for no good reason.) Having now recognized that it follows from P, if she doesn't then know Q, that would surely be a counterexample to closure. But Hawthorne's Formulation won't count it as such, since the antecedent isn't satisfied: she doesn't *come* to believe Q, since she already believed it.

Knowledge closure, so formulated, is a synchronic closure principle, imposing a condition on the set of propositions known by an agent at a time.³¹

This is, I think, the most defensible formulation of knowledge closure. Unfortunately, however, there are prominent views in the literature that are universally taken to preserve closure but are incompatible with KC. So KC – and, indeed, any principle that identifies knowledge itself as the epistemic property that is closed over inference – still doesn't suffice as the target principle for the closure debate.

However, a closely related closure principle, one which concerns *warrant* rather than knowledge, avoids the problem. To see the problem with KC, we need to introduce two distinctions: between closure and transmission, and between knowledge and warrant. The next section will do so, and will provide definitions of warrant transmission and warrant closure. With those definitions in hand, I will describe the problem with KC in §1.7. The upshot is that the knowledge closure debate really concerns warrant closure (and transmission).

1.6 Transmission and Warrant

Closure is importantly distinct from transmission. Knowledge transmission takes place when knowledge of the conclusion is *acquired* as a result of S's recognition that P implies Q. But it is compatible with transmission failure – S can't acquire knowledge of Q by recognition that it follows from P, even if she knows P – that S inevitably ends up knowing Q anyway, so that closure is preserved. Clause (d) of

³¹ For the sake of convenience, I will often speak hereafter of S's inferring from P to Q. But doing so is only shorthand for S's believing P, recognizing that P implies Q, and believing Q on the basis of that belief and recognition.

KC doesn't specify how S comes by her knowledge of Q, and so doesn't require that she acquires it in virtue of her recognition that it follows from P.

Crispin Wright has long argued that there are cases of this sort; indeed, some of his examples include the very cases that have long been cited as closure failures, such as Dretske's famous zebra case.³² While there is considerable disagreement with respect to the source of transmission failure, many closure advocates nevertheless agree with Wright that transmission does fail, and that its doing so helps explain why the inferences in zebra-style cases seem to go awry.³³

Transmission-failure cases (if they exist) are more perspicuously represented by reflecting on what it is that is preserved when transmission succeeds. S's recognition that P implies Q is obviously not responsible for the truth of the conclusion when knowledge transmits. Since the inference is valid and the premise true, the conclusion is true whether or not S recognizes the implication. Nor does such recognition secure knowledge because it secures belief in the conclusion or because it ensures that the agent believes the conclusion on the basis of that recognition. Clause (c) of KC is included precisely because an agent could recognize the implication and, nevertheless, fail to believe the conclusion (or fail to believe it in response to that recognition).

³² See the references to Wright's work in [Chapter 6, fn. 5](#). In fact, Wright is concerned with the transmission of warrant rather than knowledge and, indeed, with the transmission of the legitimacy of claims to warrant rather than warrant itself. I will explore Wright's view in [Chapters 6](#) and [10](#), and *passim*.

³³ I will discuss transmission failure and its explanation in [Chapters 4](#) and [5](#).

The epistemic property transmitted is, rather, warrant in Alvin Plantinga's sense: that which makes for the difference between mere true belief and knowledge.³⁴ This implies no theory of warrant. Warrant may or may not imply belief or truth, and it may be external or internal or have elements of both.³⁵ Plantinga-warrant – *P-warrant*, hereafter – is merely a placeholder for whatever it is that differentiates a truly believed proposition from one that is known.

I can't acquire new knowledge of Q by inference when I know it already.³⁶ I can, however, acquire a new warrant for Q that way, even if I already possess a warrant for it. Warrants are essentially ways of knowing, and I might well know Q in more than one way. The question whether transmission fails is the question whether there can be cases wherein S doesn't acquire an additional warrant for Q in virtue of her recognition that it follows from her warranted belief in P, whether or not she had a warrant for Q to start with.

If she does inevitably acquire such a warrant, then the following is true:

Warrant Transmission (WT)

Necessarily, for every agent S and propositions P and Q: if

- (a) S's belief that P is warranted while (b) S recognizes that Q follows from P, then (c) S acquires a warrant for Q in virtue of (a) and (b).

³⁴ Plantinga 1993a and 1993b.

³⁵ In §§8.5–8.6 I will argue that warrant in Plantinga's sense does, in fact, imply truth. Even if so, however, this is not a consequence of its definition alone.

³⁶ Unless I somehow lose that knowledge in the interim.

If S then believes Q on the basis of that warrant, then not only does S *have* a warrant for Q, but her belief in Q *is* warranted. The former only requires that S has such a warrant available to her, but does not imply that it is exploited as a basis for her belief; the latter does so imply. This is analogous to the standard distinction between propositional and doxastic justification.³⁷

Warrant closure can now be formulated in such a way that it follows from WT, as it should, by weakening the consequent:

Warrant Closure (WC)

Necessarily, for every agent S and propositions P and Q: if

(a) S's belief that P is warranted while (b) S recognizes that Q

follows from P, then (c) S has a warrant for Q.

While WT implies WC, the reverse implication doesn't hold: clause (c) in WC does not imply that S's warrant for Q is a result of her recognition that Q follows from P as required by clause (c) of WT. If Wright is correct, then there are cases in which S does not acquire a warrant in virtue of that recognition, so that WT is false, although she does have a warrant for Q from some other source, so that WC remains true.

³⁷ (c) should not read "S has a warranted belief in Q" because she might not believe Q at all, or only believe it as a result of wishful thinking rather than because it follows from P. Nor, I think, should (a) read "S has a warrant for P." Notwithstanding her having such a warrant, she might nevertheless believe P only as a result of wishful thinking. It strikes me as far less plausible that she must end up with a warrant for – a way to know – Q that, she recognizes, follows from P that she only believes as a result of wishful thinking.

1.7 The Problem with KC

Not only are closure and transmission more perspicuously characterized in terms of warrant; doing so is unavoidable. For it turns out that KC – unlike WC – doesn't allow for a view like Wright's. But whatever one might think of such a view, it surely should not be foreclosed by the definition of closure itself.

To see this, recall that KC reads as follows:

Necessarily, for every agent S and propositions P and Q: if

(a) S knows that P while (b) recognizing that P implies Q, and (c) believes Q on the basis of her belief that P and recognition that P implies Q, then (d) S knows that Q.

Notice that (c) specifies the basis of S's belief: she believes Q because P implies it. Suppose that she only believes Q for this reason. Q is also true, since she knows P – which requires that P is true – and P implies Q. Of course, if WC is false, the antecedent of KC doesn't ensure that she has a warrant for Q as needed in order for her to know Q; KC requires that WC is true. So suppose that WC is true. Then, since she is warranted in believing P (as required by her knowing P), and recognizes that P implies Q, she also has a warrant for her true belief in Q.

But that doesn't imply that she knows Q. Knowledge surely requires, not merely that one have a warrant for one's true belief, but that one's belief is warranted, that is, it is based on the warrant one has for it. Otherwise, one could count as knowing a proposition that one believed only on the basis of, for example, wishful thinking.³⁸ If transmission fails in this case, S doesn't acquire a

³⁸ In a similar vein, it is widely recognized that one's belief is not justified if one merely has a justification for that belief; one's belief must also be based on

warrant for Q in virtue of her recognition that it is implied by P. WC does require that she, nevertheless, has a warrant for Q (by, therefore, non-inferential means). But, we have supposed, she doesn't believe Q on the basis of that warrant; she only believes Q because it is implied by P. So she doesn't believe Q on the basis of any warrant she has; her belief isn't warranted. So she doesn't know Q. KC fails.

KC is only preserved if WT is true, that is, if recognition that P implies Q *always* delivers a warrant. Then clause (c) of KC will ensure that her belief is based on that warrant. She will then have a true, warranted belief in Q, and, therefore, know it. But to insist that WT is true just is to deny that transmission ever fails. So KC is true only if transmission never fails. And that rules out a position like Wright's – according to which transmission occasionally fails, although closure succeeds – by fiat.

KC would still be viable if Wright's view were understood to require, not only that S must *have* an independent warrant for Q when she is warranted in believing P and transmission fails, but also that her belief in Q *is* warranted, that is, that it is based on a warrant that she has. Since Q is true (because P is true, as per clause (a) of KC), she would have a warranted, true belief in Q and so know it, as required by KC.

But this additional condition should not be foisted on Wright. It would make the position much less plausible, since it would require that S believes Q. It is utterly implausible that real agents have so much as contemplated every Dretske-style Q proposition following from their ordinary beliefs; they may well

a justification one has for it. That's the point behind the distinction between propositional and doxastic justification. The same point surely applies to knowledge.

not even have the conceptual resources to do so. There being no obvious reason to think that Wright's position needs this additional condition, it would be unfair to him – and to the closure advocate in general who wishes to reconcile closure with the concession that transmission fails in some cases – to impose it.³⁹

At this dialectical stage, at least, Wright's view should not be foreclosed by the definition of closure itself. So the scenarios that view conceives of as possible – wherein, although transmission fails, S has a warrant for Q – should not count against closure. The only way to ensure this is to characterize closure in terms of warrant rather than knowledge, and so by WC rather than by KC.⁴⁰

1.8 Transmission versus Penetration

Transmission is sometimes characterized as requiring that the warrant for the premise is itself carried through the inference to the conclusion, so that the very same warrant for the premise becomes a warrant for the conclusion. Dretske, for example, so characterizes transmission in Dretske (2005, 15). But this is stronger than transmission – at least, it is stronger than I intend the term here. In Dretske (1970) he argued that a reason for believing P does not necessarily constitute a reason for believing Q: I might see that there is wine in the bottle without seeing that there isn't colored water in the bottle, although the former implies the latter. He called these "penetration" failures, claiming that they

³⁹ See also Lockhart 2018, §3.

⁴⁰ Silins 2005, §5.2 presents a similar argument. This is not a problem for KC in particular; it applies to any principle that identifies knowledge as the epistemic property closed over inference (and so also applies to the Classical Formulation as well as to Hawthorne's Formulation).

demonstrate that knowledge does not inevitably transmit, and so that closure fails.

But, as Klein (1995), Luper (2006) and others have pointed out, S's warrant for Q is grounded in her recognition that P implies Q, which is not how her warrant for P was acquired. So their warrants are distinct; failure of penetration in Dretske's sense is, therefore, unsurprising, and does not imply that S does not acquire a warrant for Q in virtue of her recognition that P implies Q.

One might think that S's warrant for P must, nevertheless, be available to S as a warrant for Q, at least when transmission succeeds. After all, Q is logically weaker (or, at least, no stronger) than P. If her warrant for P doesn't suffice for Q, and yet a warrant is generated by S's recognition that Q follows from P, then that recognition would have the seemingly magical effect of creating a new warrant for the weaker Q that was not there already in the warrant S had for the stronger P.⁴¹

Some also cite the fact that any evidence that makes P probable must render Q at least as probable.⁴² If warrant is (at least in part) a matter of probabilistic evidential support, then this would appear to suggest, not only that S's belief in Q is warranted, but that it is warranted directly by the same evidence

⁴¹ Tony Brueckner, for example, characterizes this configuration of claims as "extremely odd" (Brueckner 2000, 142, referencing Klein 1995, who endorses these claims). Both Brueckner and Klein are concerned with justification closure rather than knowledge (or warrant) closure; but the point at issue applies equally well to the latter.

⁴² Klein 1995 so argues, for example, on p. 219. This is the one positive argument for closure referenced in fn. 5.

that delivers P's warrant (whether or not it is also warranted in virtue of S's recognition that it follows from P).⁴³

However, it is relatively easy to come up with cases in which the warrant for P is clearly not a warrant for Q, and yet a warrant for Q is generated by inferring it from P.⁴⁴ Presumably I can know – and so acquire a warrant for – “the liquid in the cup is water” on the basis of how it looks, tastes, and smells.

Knowing also that water is composed of H₂O molecules, I infer that there are H₂O molecules in the cup. Presumably I can come to know that there are H₂O molecules in the cup as a result. But I surely don't acquire a warrant for “there are H₂O molecules in the cup” on the basis of how it looks, tastes, and smells alone. Someone without background knowledge that water is H₂O – or who had that background knowledge but nevertheless failed to recognize that it, together with the liquid's being water, implies that the liquid is H₂O – would not have a warrant for “there are H₂O molecules in the cup.”⁴⁵

⁴³ Ironically, Klein offers this in support of closure in the very article – [Klein 1995](#) – to which Brueckner responds, nevertheless failing to notice that it does not sit well with his claim that the justification for P might not be available as a justification for Q.

⁴⁴ For those willing to endorse the reasoning in Dretske's wine case from “there is wine in the bottle” that one knows by seeing the wine to “there is not colored water in the bottle,” that case already stands as an example ([Dretske 2005](#), 14). For Dretske is presumably right to claim that one does not *see* that there is not colored water in the bottle by looking at it.

⁴⁵ Some might object to the mobilization of background knowledge. But such mobilization is common; in Dretske's zebra case, for example, “that's a zebra” is widely taken to imply “that's not a disguised mule,” although this depends on the background knowledge that zebras are not mules. And, even if we treat the background knowledge as part of the warrant, so that my putative warrant for “there's H₂O in the cup” is based on its looking, tasting, etc., like water in conjunction with my background knowledge that water is H₂O, I

Another example: I again know that there is water in the cup on the basis of how it looks, tastes, and smells. This implies that it's not the case that both there isn't water in the cup and the universe is expanding. But it's bizarre to suggest that I can learn the latter solely on the basis of how the liquid looks, tastes, and smells, and so without recognition that the negated conjunction is a logical consequence of what I learned on that basis (namely, that it is water). And another example: logical truths follow from everything. So "that's water" implies "it's not the case that snow is and isn't white." But the basis of my warrant for the former – that the liquid looks, tastes, and smells like water – surely does not, on its own, serve to deliver a warrant for an unrelated logical truth.

Moreover, if the argument for closure is based on the claim that one's warrant for P is available as a warrant for Q – as the appeals to logical strength and probability suggest – then there is no need for S to recognize that Q follows from P at all. That recognition, therefore, plays no role in ensuring that S has a warrant for Q. The closure principle this supports is, then, as follows: if S's belief in P is warranted and P implies Q, then S has a warrant for Q. But this is obviously false. If S doesn't even recognize that P implies Q, then she certainly doesn't inevitably end up with a warrant for Q, even if she has a warranted belief in P. The intuitive motivation behind closure, after all – that S's *recognition* that P implies Q allows for S's acquisition of a warrant for Q – is now irrelevant.

obviously don't need that background knowledge in order to be warranted in believing that there's water in the cup. So what I do need for the latter warrant doesn't suffice on its own to warrant "there's H₂O in the cup." That background belief is only relevant to my warrant for "there's H₂O in the cup" in virtue of its role in facilitating the inference from "there's water in the cup" to that proposition.

Indeed, one might as well just point out that, since P implies Q, anything that makes P true makes Q true as well, and offer *that* as an argument for closure. It supports the same untenable version of closure, after all. And while it is disputable whether knowledge of P requires that P be probable on one's evidence, it is beyond dispute that knowledge requires truth.⁴⁶ But an argument for closure that appeals solely to the fact that valid inference is truth-preserving is no argument at all (or is, at best, a question-begging one).

So transmission, as I will use the term, only requires that S acquire a warrant for Q as a result of inference from P; it does not require that S's warrant for P itself suffices as a warrant for Q.⁴⁷ This does mean that Dretske's penetration-failure argument for closure failure doesn't succeed; the argument of this book does not appeal to that argument. It also, however, undermines appeal to logical strength or probability in defense of closure.

[Chapter 2](#) presents another of Dretske's arguments, one which appeals to *prima facie* counterexamples to closure. That argument will structure the examination of the options available to the closure advocate in [Chapters 3–9](#).

⁴⁶ See [Williamson 2014](#).

⁴⁷ For more on the suggestion that warrant for P suffices as a warrant for Q, see [Chapter 7](#).